Commonwealth of Kentucky

**Exchange 2003 Global Address List Synchronization
Design and Implementation**

Version 1.4
1/2/2008

The purpose of this document is to provide an interested Commonwealth of Kentucky (COK) government entity that is compliant with all stated qualifications listed in this document, the necessary information to procure, install and configure the necessary hardware and software to successfully automate the synchronization of their Exchange 2003 Global Address List with other participating COK entities and to achieve a fully automated solution for keeping all participating GAL's across the Commonwealth up-to-date.

Table of Contents

Project Background

Most, if not all, business entities within the Commonwealth of Kentucky (COK) utilize Microsoft Exchange 2003 as their enterprise e-mail platform. Most entities within the Executive branch of Government utilize the services of the Commonwealth Office of Technology's Exchange implementation for their e-mail needs, while members of the Legislative, Judicial and Department of Education all implement and manage their own instance of Exchange (their own Organization) within their own Active Directory forest.

Most branches of Government implement separately due to the foundations upon which separation of Government principles are founded on and to provide an ultimate boundary for the purposes of data security. Although this separation brings with it many operational advantages, one of the issues that arise as a by-product of multiple instances of Exchange in different Organizations is that there are multiple Global Address Lists (GAL's) that are not automatically synchronized and kept up-to-date with changes in Commonwealth personnel.

Under Exchange 5.5, the synchronization of GAL data was performed by a series of manual export / data massage / import routines that were left to each entity to perform for themselves at their own discretion. These imports were cumbersome, time-intensive, and often prone to errors in either operation or data validity. Factor in that each agency performed their own export/import cycle on their own schedule, and the problem of data concurrency and validity becomes an issue.

Once all entities moved to Active Directory and Exchange 2003, the problem, as stated above, was not resolved via new features of either product. All entities moved to offering their GAL data in updated formats so that other entities could 'swap' GAL information under the new environment and Since Exchange 2003 now relies solely on Active Directory for its directory data, updating the GAL became an exercise in updating Active Directory. Unlike Exchange 5.5, the tools for doing so are not as intuitive to use and the data formats used are not as easily (and uniformly) manipulated as they were in Exchange 5.5.

Although there have been some scripted import/export processes put in place in different entities, there is no current enterprise approach to solving the GAL synchronization problem in a manner that is consistent, timely, automated and concise. The goal of this project is to address these concerns with a solution that provides for an enterprise approach to synchronizing the Global Address Lists across the Commonwealth of Kentucky.

Business Value

Business valuation on a solution such as this is probably best measured by the impact within each entity by them not having to manage this process (and data) going forward and the Global Address Lists across each entity staying consistent and uniform. Once each entity enrolls in the 'GALSync' process, cross-entity messaging within the Commonwealth will enjoy a much higher level of address list concurrency with personnel changes (additions, changes, deletions from the entity GAL) being reflected in the other entities GAL within days rather than months.

The other benefit from this is that each entity will be able to control what data is published to each 'subscribing' entity and will be able to exercise a fine level of control over how their publishing data appears to other entities. This also lends itself to standardization and provides a way for each entity to continue to publish their respective GAL data without having to worry about individual processes for importing/exporting with other entities due to messaging system versioning differences within the Microsoft platform stack.

Solution Scope

The scope of this solution is bounded by the fact that enterprise GALSync is implemented via a localized implementation of a 'GALSync' server that synchronizes your entity's data with a centralized (nonpartisan) instance of Active Directory that serves as a repository for other entities to read (but not update) your entity's data while your entity is able to read (but not update) other entity data. It is your GALSync server that does the work of updating the GAL with the data in this nonpartisan Active Directory forest and the configuration allows for the synchronization of as little (or as much) data from other entities to your GAL.

COT maintains the nonpartisan forest hardware and performs the operational duties to keep the forest functional, available and protected from disaster (performing backups, patching, etc.) and is a user of the solution in much the same way as everyone else.

This document details the overall structure of the enterprise solution and details the necessary hardware and software components that are necessary to implement this solution. The setup and configuration are also detailed.

Note

Any custom business requirements that your entity may have cannot be covered via a general setup and configuration of the GALSync software, and those advanced configurations are left to you to configure. This document explains how to perform 'whole-environment' GALSync with other entities within the Commonwealth and if there is a need to customize this behavior then advanced configuration (which is outside the scope of this document) will be required.

Business Requirements

This solution was comprised for all participating state entities, with special input from the Commonwealth Office of Technology (COT), the Kentucky Department of Education (KDE) and the Legislative Research Commission (LRC). All three parties contributed their own unique business requirements to the overall solution and were very gracious in participation on the project to help ensure that the as-delivered solution was flexible enough to accommodate some very advanced business requirements that needed to be represented in the synchronization process.

The primary business requirement, being common between all of the aforementioned entities, is being able to exchange e-mail-related addressing information with other COK entities on a regular basis in an automated fashion that allows their own address lists to stay current.

Specific business requirements that each entity has described as being business-critical are noted below:

Commonwealth Office of Technology (COT)

- Distribution list membership cannot be altered (user removed/re-inserted) when a modification to a member of the list is processed.

Kentucky Department of Education (KDE)

- Requires the ability to filter specific recipient types (namely students from school districts) from synchronization.
- Distribution list membership cannot be altered (user removed/re-inserted) when a modification to a member of the list is processed.
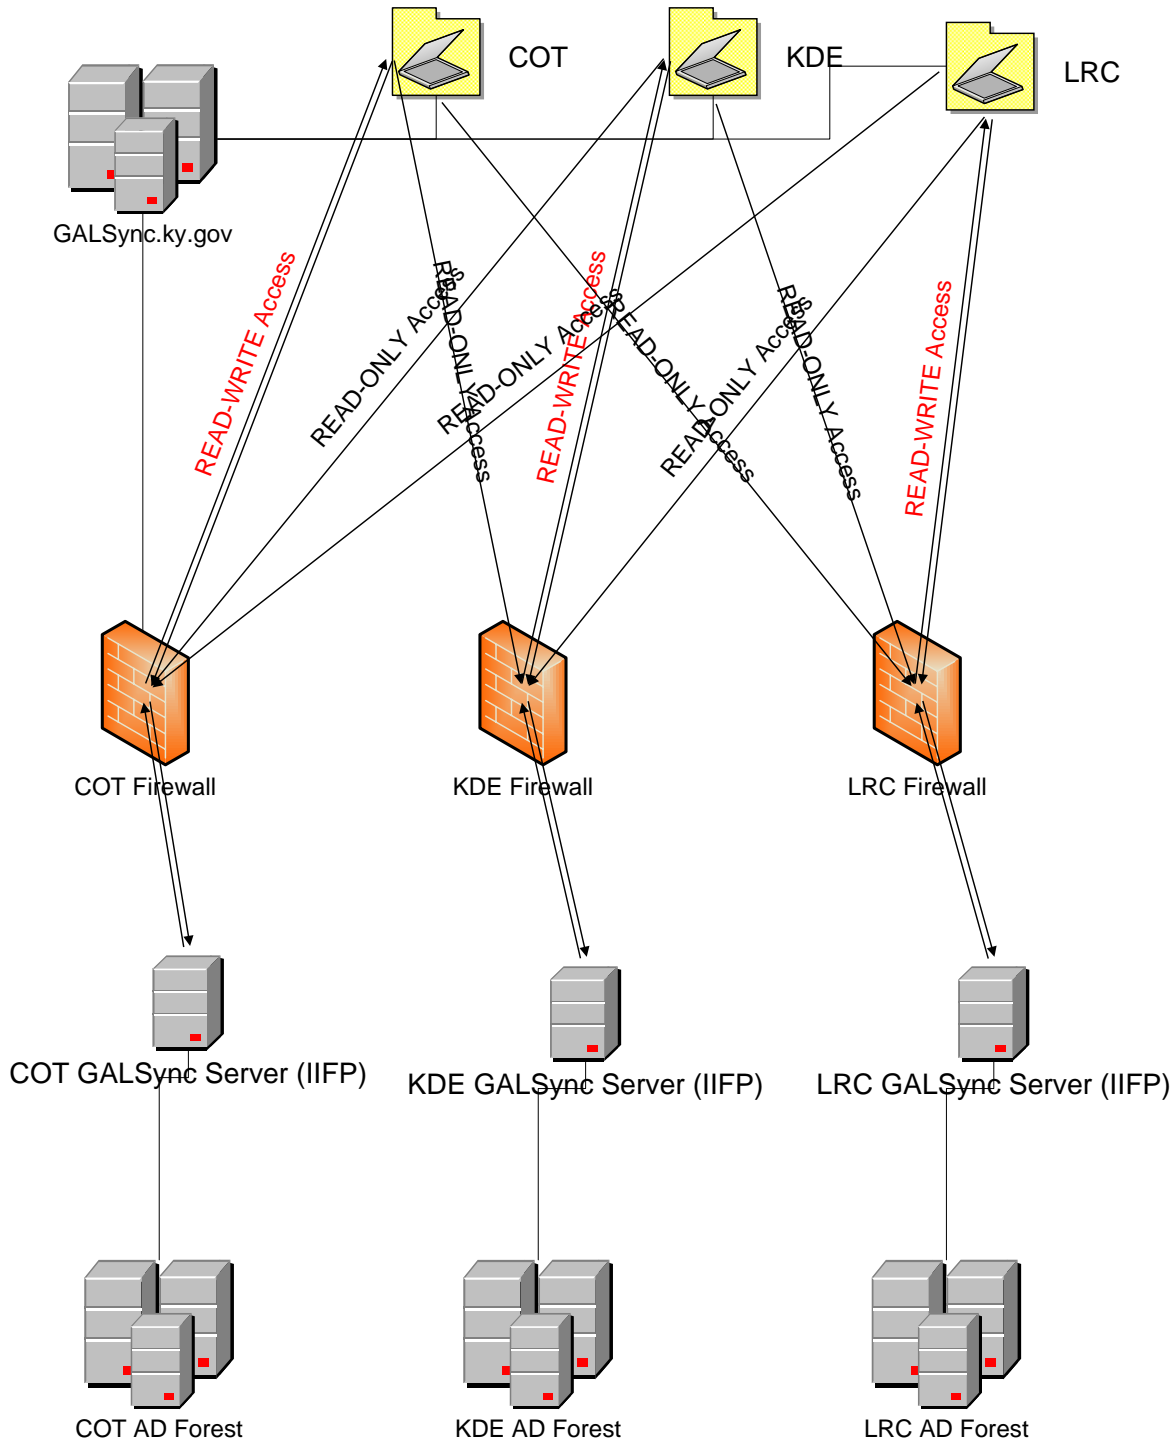
Legislative Research Commission (LRC)

- Distribution list membership cannot be altered (user removed/re-inserted) when a modification to a member of the list is processed.

Additional Entities

- To be determined (and configured) by the entity themselves.

Solution Architecture

The overall solution architecture is based around a 'subscriber/publisher' model that each participating entity follows. The depiction of the logical architecture is shown below:

The Executive Branch's Commonwealth Office of Technology (COT) will provide and maintain the GALSync forest. This forest will only contain the synchronized contacts of the participating entities and should not contain any sensitive information from any participating entity (as long as the synchronization configuration guidelines are followed).

Each participating entity will have their own IIFP/SQL server that serves as their 'GALSync' server. This server will make an outbound connection through their firewall, and connect to the 'GALSync' forest that resides in the COT DMZ (E-Gov) zone.

Each entity will be provided with a GALSync forest service account that will have read-write permissions on the entity OU where all mailbox-enabled users are written to and will have read-only access to all non-entity OU's (other governmental participants) where the GALSync server will be configured to read in all mailbox-enabled users that have been exported from other entities. The advantage of this method is that each entity is capable of configuring who they want to see in their own GAL and can add/remove other entities at will. Each entity will also have only enough permissions to update *their own* data in the GALSync forest and no permissions to update anyone else's data.

Due to the fact that all GALSync data resides in a common and secured forest, there is no need for any participating entity to allow connectivity from any other entity through their own firewall, as all connections for GALSync are originated from within the participating entity outbound to a common location, and no inbound connections are required or should be configured specifically for this solution.

This also allows an entity to participate when they want, and how often they want (recommended guidelines are listed in this document) without impacting any other participating entity.

Note

Although COT, LRC and KDE are commonly-referenced in this documentation, the solution is architected such that any approved agency can participate in the GALSync process. This process is not unique to the agencies listed above, and they are used for representative purposes only.

The 'subscriber/publisher' model is simply repeated once per additional entity that participates in GALSync.

Solution Requirements and Responsibilities

This section details what will be required, from an entity perspective, to quality for the GALSync service and install, run and manage the GALSync server.

Responsibilities

- **COT Responsibilities** - COT will be responsible for the implementation of the first IIFP server and staging forest that all participants will use to house their contacts which includes any hardware and software costs. COT will maintain the hardware and software to insure proper functionality of the GALSync forest's domain controllers and will make every reasonable effort to ensure the continuous availability of the services they provide.
- **Participants Responsibilities** - All participants understand that they are responsible for acquiring their own GALSync server and any performing any configuration that will be required to successfully participate in the GALSync program. The entity will be responsible for any support required to implement this solution as well as any on-going support that may be required to ensure the proper long-term functionality of their GALSync server. This includes, but is not limited to the day-to-day server operations, backup and restoration of the IIFP database, GALSync server configuration, backup and restoration of the entity Active Directory database before performing the initial synchronizations. COT will not be responsible for the validity of data in each entities forest, either on the import, synchronization or export of said data to and/or from the GALSync forest.

Qualifications

The qualifications to participate in GALSync are as follows:

- An entity's GAL is their property and the use of it by someone else needs to be approved by the proper management of that entity.
- An entity needs to be part of the Commonwealth's network which has been referred to as "The Kentucky Information Highway" per KRS 45A.605.
- Completion and submittal of the "GAL Synchronization Letter of Acknowledgement" (included in Appendix A of this document) to COT for approval and GALSync forest service account creation
- An identified entity contact who will serve as the 'point contact' for all items related to GALSync for that entity between themselves and COT.
- The ability to legally dispose of the following information about all users who will be synchronized to other entities via the GALSync process:
  - First Name
  - Last Name

- o Display Name
- o E-mail alias
- o Work Address
- o Work City
- o Work State
- o Work Zip Code
- o Title
- o Company
- o Department
- o Office
- o Work Phone
- o Work E-mail Address

Environmental

- Active Directory running in Windows 2003 forest and functional mode is a recommendation, but not a requirement.
- Exchange 2003 schema extensions and Exchange 2003 operational (prior versions of Exchange are not supported for GALSync)
- One external IP address (static) on the external interface of the entity firewall that the GALSync server will always 'appear' to be coming from.

Hardware

The GALSync server is to be a single, stand-alone server. No components of the solution are to be clustered or spread across multiple servers. Due to the nature of the synchronization schedule, a well-equipped virtual machine may be the proper (and economical) choice for many entities to implement this solution on.

The server should also have a gigabit/GB (recommended) or 100mb/sec (minimum) network connectivity. If possible, locate the GALSync server as close (subnet-wise) to the most Active Directory domain controllers it will be communicating with so as to minimize router hops and maximize network communications.

Since the GALSync server resides on the entity internal network, IP addressing and other related requirements are at the responsibility of the entity.

Although this documentation recommends a physical server for performance and supportability, a virtual machine is also a very viable option. Recommended configurations for both are given below:

Physical Hardware

- Server-class hardware with 2 or more Pentium 4 Xeon CPU's, 3Ghz or higher (2 or more dual-core CPU's recommended, 2 Ghz or higher if possible)
- 4GB RAM (more if possible)
- Hardware RAID controller with the following spindle arrangement:
    - 1 RAID-1 array dedicated for the operating system, IIFP and SQL Server binaries, temp directories and page file.
        - This should be the 'C:' drive
    - 1 RAID-1 array dedicated for the SQL Server Transaction logs
        - This should be the 'D:' drive
        - Ensure that the stripe size is 64K or larger
    - 1 RAID-1 or RAID-5 array dedicated for the SQL Server database files
        - This should be the 'E:' drive
        - Ensure that the stripe size is 64K or larger
    - 1 RAID-1 or RAID-5 array dedicated to database backups and the tempdb database/logs
        - This should be the 'F:' drive
    - Assign the CD-ROM/optical drive to be 'Z:'

Note

Each RAID array should be serviced by dedicated spindles, if possible. Combining spindles will create I/O contention and will lead to degraded performance

Virtual Machine

- Configure dual-CPU support in the VM, if possible
- Configure 4GB of RAM (or more, if possible)
- Create multiple, dedicated virtual hard disk files for each of the 4 volumes, described above, that will be needed by the VM. Try to locate these hard drive files on separate physical spindles on the host machine, if possible, to reduce I/O contention. If the VM is being back-ended by a SAN, then most I/O problems will likely be minimized by the SAN controller(s), but the same logical layout (and dedicated SAN disks, if possible) should be followed for consistency's sake.

Note

Note that the IIFP will work fine in a virtual machine, but is not currently "supported" by Microsoft Product Support Services (PSS). If you have a need to call PSS in regards to an operational issue with your IIFP server, then know ahead of time that PSS may "make" you move the IIFP to a physical server before they offer support.

Note

If you have more than 4GB of RAM available to the physical machine or virtual machine, SQL can be configured to take advantage of it via PAE and AWE, so don't let 4GB be an artificial ceiling for RAM for the IIFP server. Although Windows and IIFP won't make direct use of the memory, the IIFP will indirectly benefit from the memory due to SQL being able to take advantage of the larger memory set. Note that there are additional configuration steps that need to be taken in order for Windows and SQL to use this memory over 4GB, but those configuration steps will be included in the final documentation.

Software

- Microsoft Windows Server 2003 (Enterprise Edition) with SP2 and all applicable hotfixes **(32-bit edition)**
- Microsoft .NET Framework 2.0 and post-RTM hotfixes
- Microsoft SQL Server **(32-bit edition)**
  - 2000, Enterprise Edition, Standard Edition, or Developer Edition, with Service Pack 3a (SP3a) or higher (SP4 recommended)
  - 2005, Enterprise Edition, Standard Edition, or Developer Edition with Service Pack 1 or higher (SP2 recommended)
- Microsoft Identity Integration Feature Pack (with SP2 integrated) – available at: http://www.microsoft.com/downloads/details.aspx?familyid=d9143610-c04d-41c4-b7ea-6f56819769d5&displaylang=en

Operational

The intent of the solution is to be 'automatic' once it is fully configured (per this document) and verified to be working correctly. That being said, there will still be a need for someone within the agency to 'own' the GALSync process and the management of the server (patching, availability, etc.). the GALSync server should not require much, if any, day-to-day care and feeding, but the service owner should periodically check the server to verify proper operation and resolve any errors that may be outstanding and to perform any required manual synchronizations in due course of troubleshooting, etc.

The GALSync server owner will have the ability to delete any and all contacts in their own forest that are synchronized from the GALSync forest (other agencies) but will not have the ability to delete their own contacts unless they coordinate with COT and/or use a tool (outside of the IIFP) to do so (ldp, etc.), so plan your first synchronization carefully because the cleanup of such will be tedious!

GALSync Server Installation and Configuration

This section details the steps necessary to install and configure the appropriate software to ensure that your agency GALSync server operates and interacts with the other agency information correctly. Step-by-step procedures are included where deemed necessary (and outside the scope of 'normal' OOB (out-of-box) configurations).

This section also assumes that you have the hardware (or virtual machine) configured appropriately, as per the previous section.

Note

Throughout this section of the document, the phrase 'complex password' is used to denote a password that meets the following requirements:

A Minimum length of 8 characters
Comprised of upper and lower-case alphabetic characters (A-Z; a-z)
Comprised of numeric characters (0-9)
Comprised of special characters (punctuation, spaces, symbols, etc.)

Server Operating System

1. Install Windows 2003 R2 Enterprise edition (32-bit) onto the C: volume and accept all defaults during the installation. If you are not using a 'SP2-integrated' version, then install Windows 2003 SP2 immediately after the server completes the build process.
    a. Install any additional components (SNMP, etc.) that may be needed for your environment.
2. Install the .NET Framework 2.0
3. Install all applicable hotfixes (be sure to also install the .NET Framework 1.1 hotfixes, along with the .NET Framework 2.0 hotfixes)
4. Ensure that the network connection is operating at the highest possible frame rate
5. Ensure that terminal services (RDP) has been enabled in remote administration mode
6. Reassign the drive letter of the CD/DVD-ROM to be Z:, then proceed to format the D:, E: and F: volumes with NTFS
7. Join this GALSync server to the root domain of your forest (or whichever domain you wish the contacts to be stored in. This text assumes that location to be the root domain of the forest).

SQL Server

Note

Perform the following while being logged on as a server administrator.

---

1. Insert the SQL 2005 Standard Edition CD (or mount an ISO image) and proceed to allow the setup program to install the native client and setup support files.
2. The setup wizard will then continue and inspect the server configuration. Click [Next] to continue through the wizard:

3. The following dialog will test the server prerequisites for installing SQL 2005. Note that although IIS isn't required for this implementation, if you wish to leverage SQL Reporting Services at some point in the future to write reports against the MIIS database, then you should stop now, install IIS, and then re-run this setup program so that the SRS component will install correctly.

   Also note that since this installation is happening in a VM environment, the hardware requirements will almost always be flagged. This is of no concern as long as you are confident that the VM is configured correctly (with enough RAM).

Click **[Next]** when ready and enter in the agency name and Product Key when prompted and click **[Next]** again to proceed to the component selection dialog.

4. Enable the components for installation, as shown below, and click **[Next].**

5. Select the 'Default Instance' as the instance name, and click **[Next].**

6. Select the option (shown below) to configure all of the SQL Services to run as the local system account and enable the SQL Server and SQL Server Agent services to be started after setup ends, and click **[Next].**

7. Choose 'Windows Authentication Mode' as the authentication mode to use for the installation, and click **[Next].**

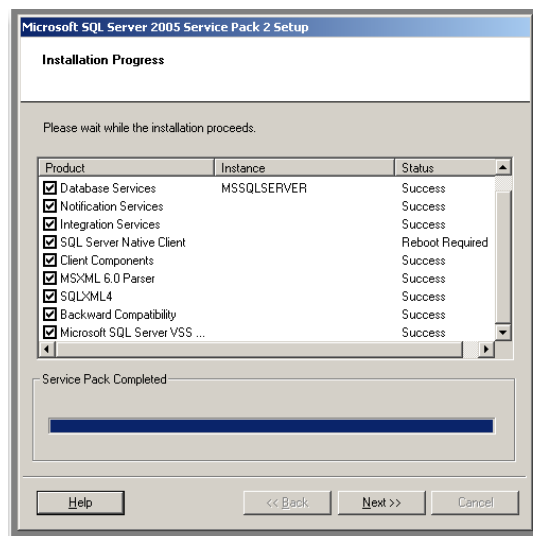8. Choose the defaults for the Collation Settings, as shown below, and click **[Next].**

9. Verify the defaults on the 'Error and Usage Report Settings' dialog. Be sure to uncheck both checkboxes, and click **[Next].**

10. Click the [Install] button on the final 'Ready to Install' dialog and allow the installer to finish the setup of the SQL instance. Be prepared to wait for a bit, as the setup process takes several minutes.

11. Once all of the components are labeled as 'setup finished' the installation is complete; click on the [Next] button

12. Click the [Finish] button in the 'Completing Microsoft SQL Server Setup' dialog once you have reviewed the installation summary (if applicable).

Once the SQL Server setup has been completed, the latest service pack for SQL 2005 (SP2, at this writing) must be installed. The following procedure details the installation of this update.

1. Locate the SP2 binaries on the COK GALSync Resources CD (or download from www.microsoft.com) in the \SQL 2005 folder. Double-click on the 'SQLServer2005SP2-KB921896-x86-ENU.exe' file and let it being extraction.

2. Once the wizard starts, click [Next] to proceed past the welcome page

3. Select the appropriate radio button to accept the license agreement, and then click [Next].

4. Ensure that all of the components are selected for update, as shown below:

and then click [Next].

5. Ensure that the appropriate authentication method is selected, as shown below:

and then click [Next].

6. Verify the defaults on the 'Error and Usage Report Settings' dialog. Be sure to uncheck both checkboxes, and click **[Next].**

7. The service pack installer will then search for locked files and owning processes and will list them in the dialog below:
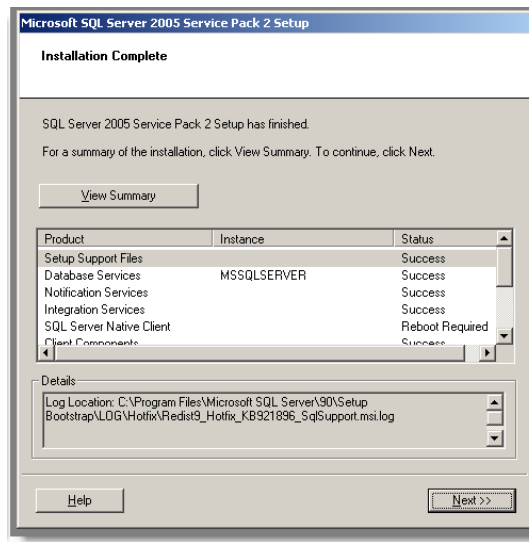

You have the option of either manually stopping all of the SQL-related services, clicking refresh and proceeding with the install, or you can alternatively click [Next] and allow the installer to go ahead and lay down the updated binaries, but a reboot will be necessary after the service pack finishes installing. Click [Next] when ready to continue.

8. Click on [Install] in the 'Ready to Install' dialog box to begin the setup process.

9. You will then see an 'Installation Process' dialog that will show how the application of the service pack is proceeding. When it is finished, the dialog should look like the
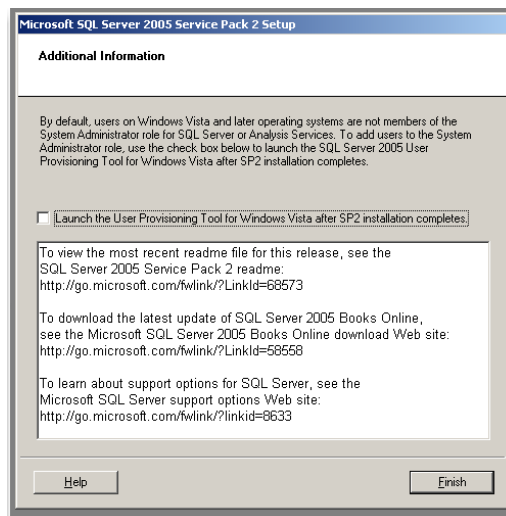

following:

Click [Next] to continue.

10. Click [Next] to continue through the installation summary dialog

11. Be sure to uncheck the checkbox that would otherwise start the user provisioning tool for Vista. This step is not necessary for the GALSync implementation:

Click [Finish] when complete.

12. Perform a restart of the server so that the files replaced during the service pack installation can be initialized upon the next boot-up.

Identity Integration Feature Pack

This section details the installation and configuration of the IIFP. Make sure you have completed the following tasks prior to performing the steps outlined in this section.

IIFP Installation Prerequisites

1. Have completed and submitted the "GALSync Letter of Acknowledgement" (see Appendix A) to COT and have received subsequent approval back to the agency contactee with approval and details on the agency-specific service account that is to be used to connect to the GALSync forest.
2. Installed and configured the Agencie's GALSync server and base OS per this document
3. Installed SQL 2005 and updated it with the latest service pack (per this document)
4. Create a domain-level service account, named 'COKGALSync' to be the service account that the IIFP runs under.
   a. **Be sure this account has a complex password!**
   b. **This account should not be a member of any administrative groups!**
   c. **Locate this account within the domain as you see fit.**
   d. **The account should have its password set to never expire and does not require change at next logon**
   e. **Do not limit the logon hours or workstations to which this account can log on from.**
   f. **This account should NOT have an Exchange 2003 mailbox.**
5. Configure this account to be a local administrator on the GALSync server
6. Configure this account to have the 'Replicate Directory Change' permission on each domain where the IIFP will have to read/write updates to/from. In practice, this is likely each and every domain that has Exchange recipients in your forest and the domain (assumed to be the forest root) where the contacts will be stored.
7. Created a 'State Contacts' top-level OU in the root domain and have assigned the domain-level GALSync service account to have 'Full Control' permissions and that they apply to 'this object and all child objects' as well.

   More information on this permission (and its requirement) can be found here:
   http://support.microsoft.com/?id=303972

   A copy of this QArticle is also included in **Appendix C - How to grant the "Replicating Directory Changes" permission for the Microsoft Metadirectory Services ADMA service account**
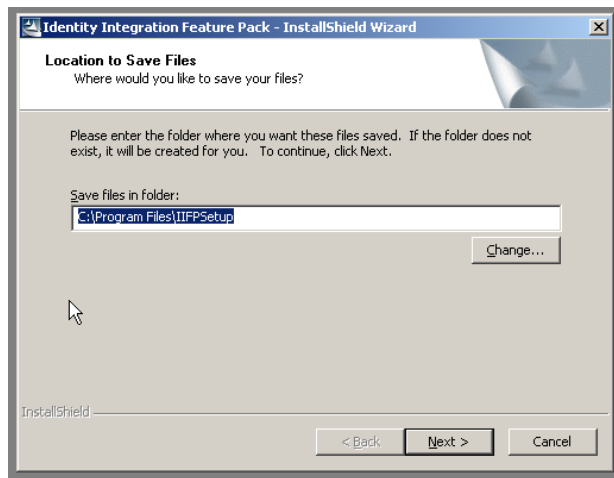
IIFP Installation and Initial Configuration

Once the aforementioned prerequisites have been met, proceed with the following steps to install and configure the IIFP on the agency GALSync server.

Browse to the 'IIFP' folder on the 'COK GALSync Resources' CD.
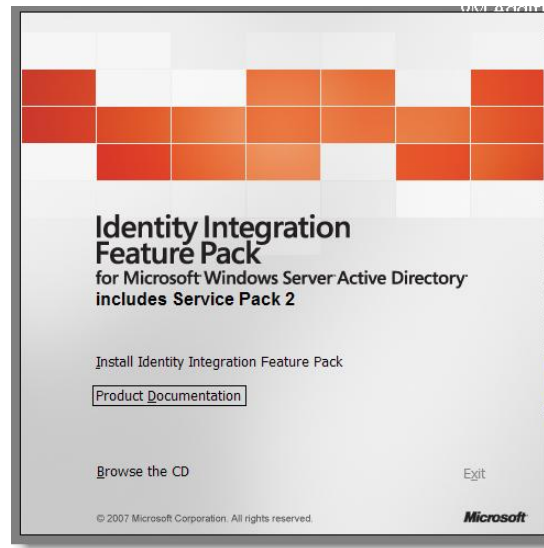
From there, execute the 'Identity Integration Feature Pack with SP2.exe' file to begin the installation of the IIFP.

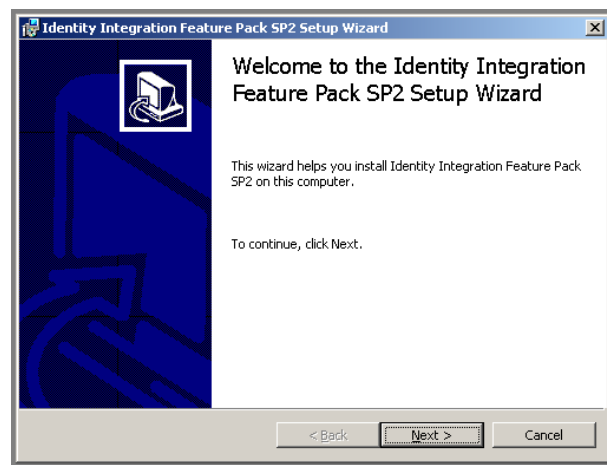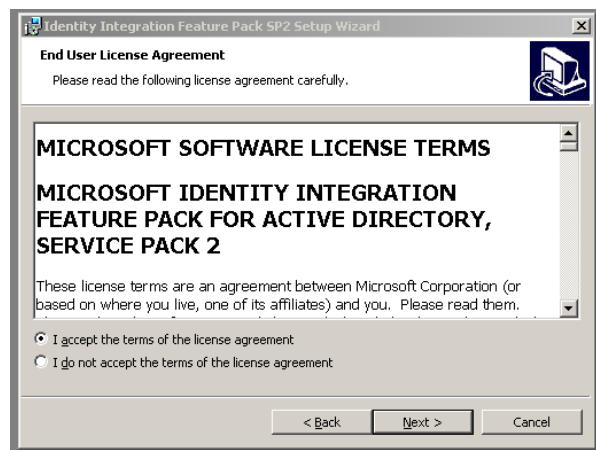8. Select the default extraction folder

and click [Next].

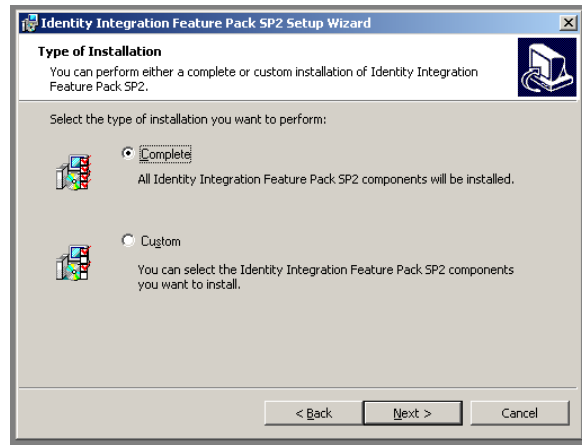9. Click on the 'Install the Microsoft Identity Integration Feature Pack' link

10. Allow the installer to load and begin extracting the installation wizard. Once the main installer UI loads, then click the [Next] button to proceed
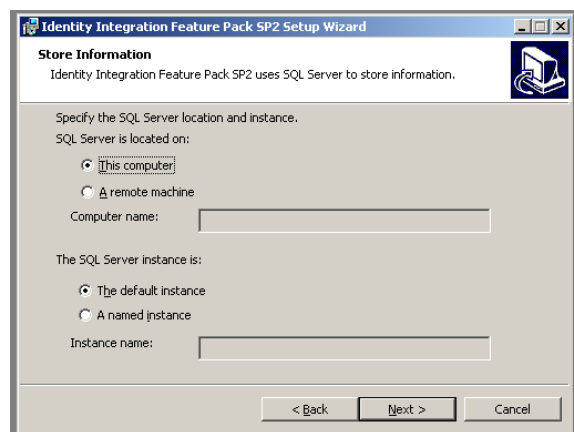

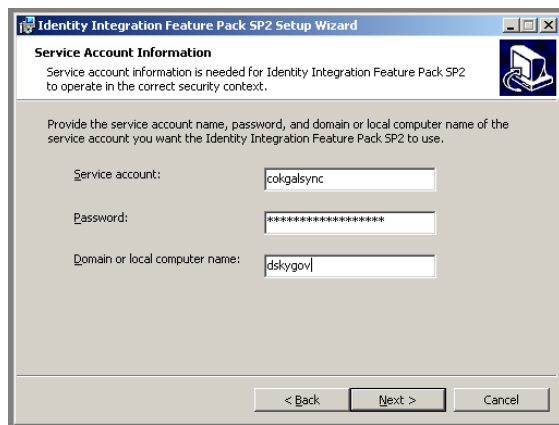11. Accept the end-user license agreement, and click [Next] to continue

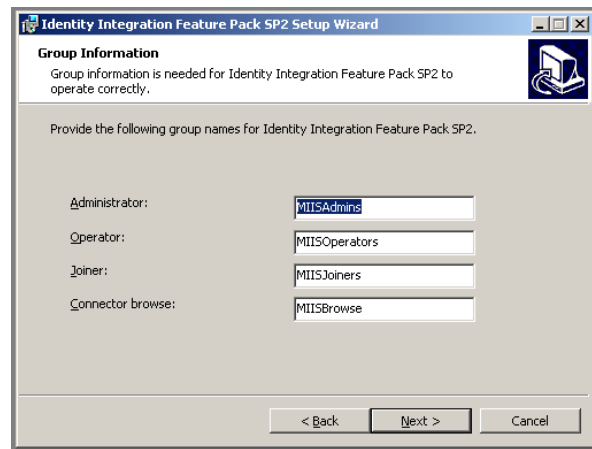12. Verify that 'Complete' is chosen as the installation option, and click [Next] to continue

13. Verify the SQL instance information and click [Next] when ready to proceed

14. Enter in the IIFP service account username, password and domain membership (be sure to use the NETBIOS name here) in the following dialog, then click [Next]
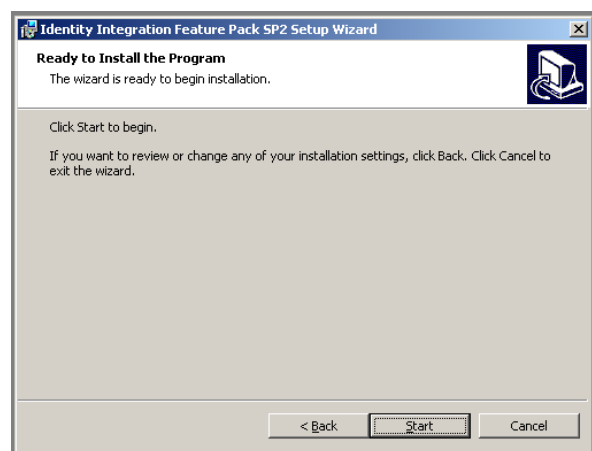
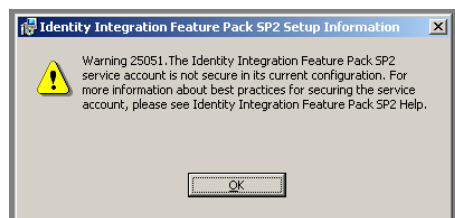15. Verify (and accept) the default group names and click [Next].

These are local security groups that will be created on the IIFP server so that you can delegate administrative permissions to others in your environment, if you wish.
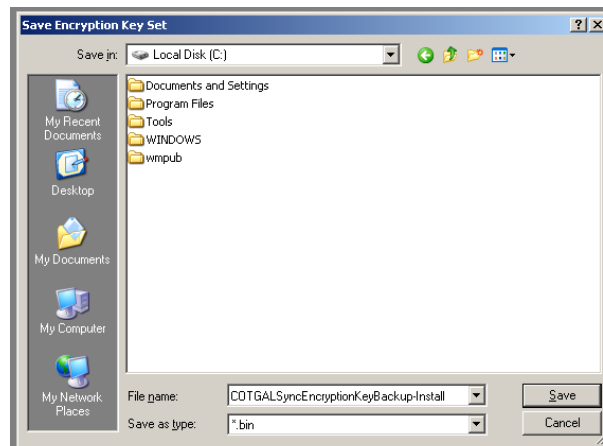
16. Click [Start] to begin the installation

17. Click [OK] to acknowledge the 'security' dialog. We will correct this once the installation is finished.
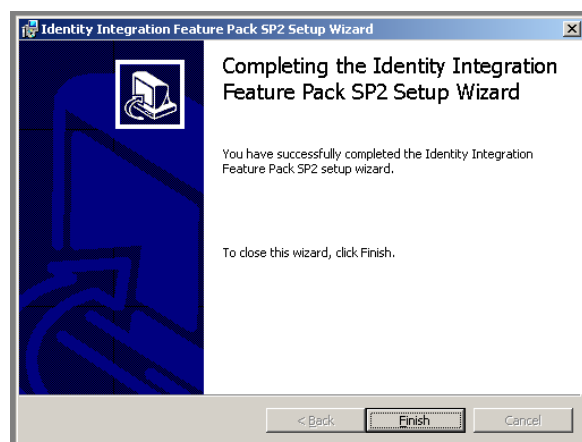
18. Click [OK] to acknowledge the security keys dialog. We will locate the backup in the next step

19. Locate the backup .BIN file somewhere easy to find (the root of the system drive is used in this example). Name the key accordingly (COTGALSyncEncryptionKeyBackup-Install.bin used in this example), then allow the wizard to continue by clicking [Save]

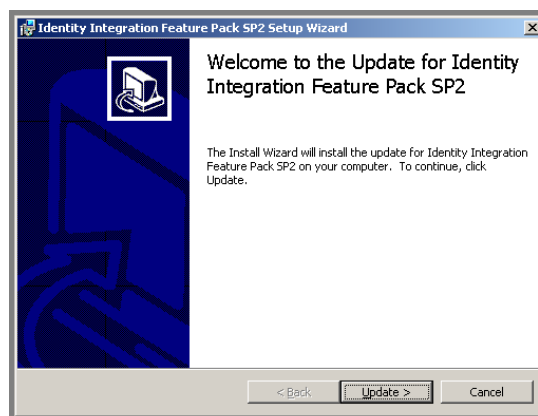20. Click [Finish] once the wizard completes.

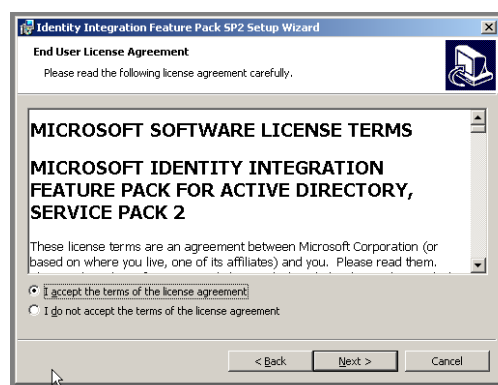21. Click [Yes] to log off, then log back on as the IIFP Service account.

22. After logging back on, open up explorer and navigate to the 'COK GALSync Resources CD' (or mount the ISO Image) and browse to the 'IIFP' folder. Execute the 'IIFP2003Update-KB938015.msp' file to begin the installation of the post-SP2 hotfix that is required for this implementation.
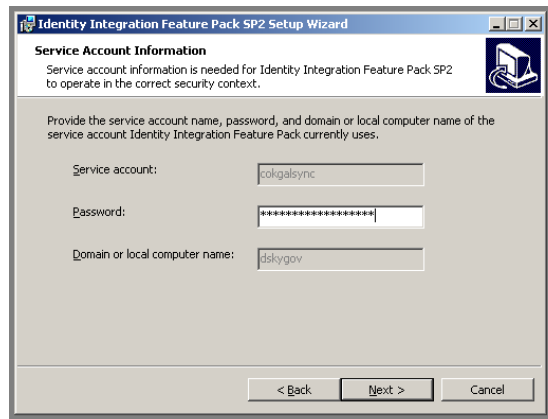
23. Click [Update] on the hotfix welcome screen and proceed through the wizard to install the hotfix.

24. Accept the terms of the licensing agreement, and click [Next]

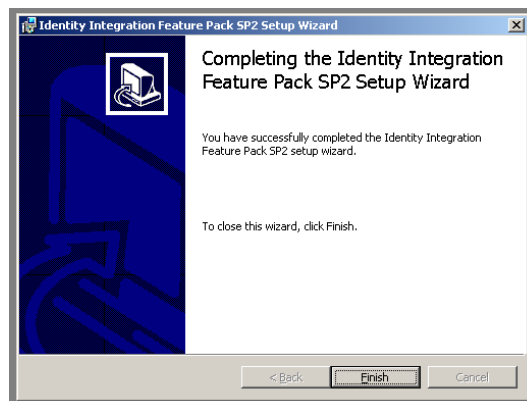25. Enter in the IIFP Service account password where indicated, and click [Next]

26. Click [Yes] to proceed with the update. Neither the database backup or key
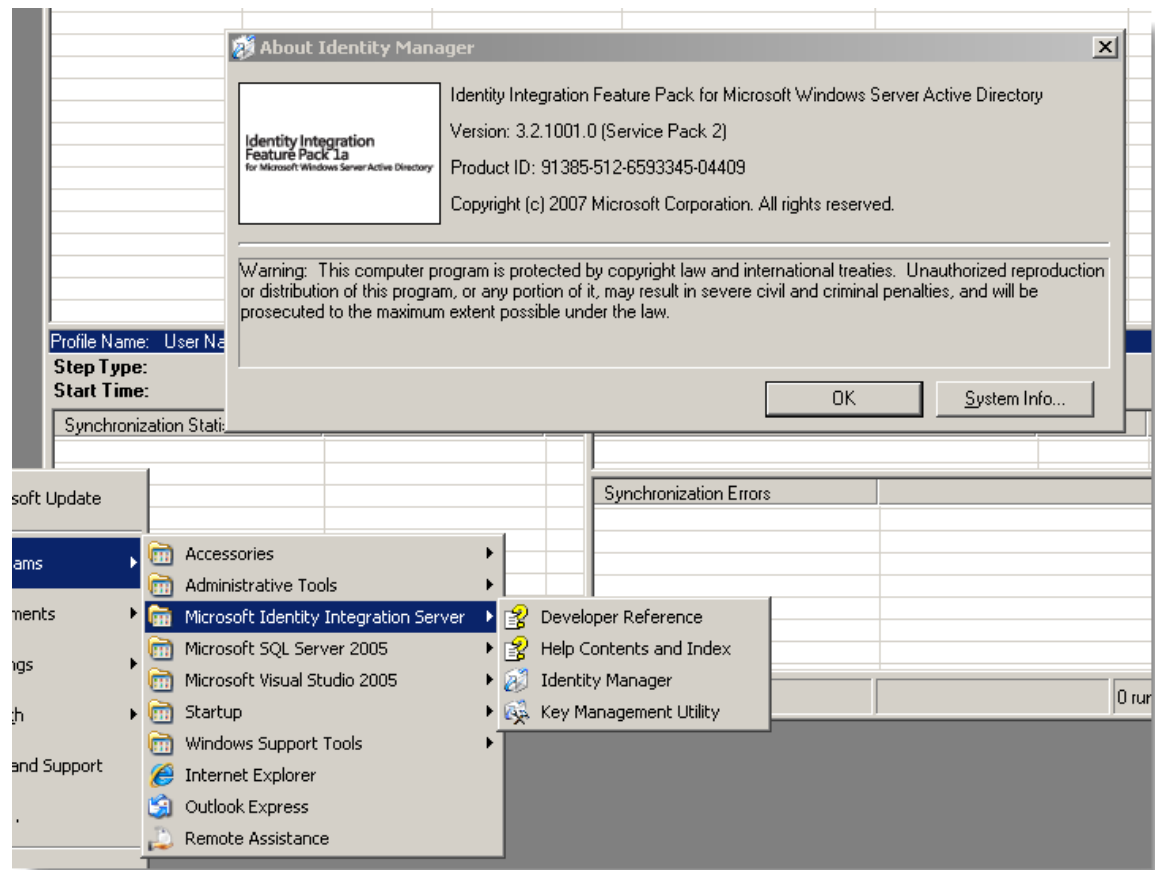

backups are necessary at this point, due to there not being any data (of value) in the database. We will back up the key set later in this section of the document.

27. Click [Finish] when the wizard completes

28. Start the Identity Manager application (as shown below) then click on Help ->
About and verify that the version number is reported as '3.2.1001.0'

This version number indicates that the post-SP2 hotfix installed correctly.
The version number for SP2 (as released) is 3.2.559.0.

29. In the next steps, we will perform a backup of the SQL Databases and the
IIFP encryption keys. These backups will then need to be stored off-server in
a protected location in the event that the server ever needs to be reinstalled
or returned to its 'before GALSync/pristine' state.

30. Open the SQL Server Management Studio, as shown below:

31. Click [Connect] to connect to the default instance on the SQL Server:

32. We will begin by making a backup of each of the databases on the server. Once the SQL Management Studio loads, drill down through the server and expand the databases node, then the 'System Databases' node, as shown below. Repeat the following procedure for each of the system databases, and then one final time for the 'MicrosoftIdentityIntegrationServer' database. The MASTER database is shown here for illustrative purposes, but all databases will back up in the same manner.

   a. Right click on the database, select 'Tasks -> Backup'

   b. Make note of the location of the backup destination (c:\program files\microsoft SQL Server\MSSQL.1\MSSQL\Backup) and the name of the backup file (this will vary for each database – just accept the

default) and then click [OK] to begin the backup process.

c. Click [OK] to acknowledge the successful backup in the dialog that pops up.

d. Repeat for the next database until finished (skip tempdb).

33. Open the folder that contains all of the backup files and copy them to a safe location that will be considered permanent storage (CD-ROM, protected file server, etc.)

34. Next, open the IIFP Key Management Utility:

35. Once the key management utility starts, verify that 'Export Key set' is chosen and click [Next]

36. Enter in the credentials requested and the domain name, and click [Next]

37. Determine an appropriate location and name for the exported key file (use something descriptive so you can identify it later, if it is ever needed). And then click [Next]

38. Click [Finish] to acknowledge the actions

39. Click [Close] to exit the wizard.

40. Open the folder that contains the backup key file and copy it to a safe location that will be considered permanent storage (CD-ROM, protected file server, etc.)

41. Now, open Computer Management and select the 'Services' node and then highlight the 'Microsoft Identity Integration Server' service, right click it and select 'Stop' from the context menu.

42. Once the service stops, re-open (or switch back to) the SQL Management Studio and then select the 'MicrosoftIdentityIntegrationServer' database, right-click on it and select 'Tasks -> Detach' from the context menu.

43. Be sure to select the 'Drop Connections' and 'Update Statistics' checkboxes, then click [OK]

Click [OK] to acknowledge the dialog that pops up warning that the database is no longer accessible.

44. Using explorer, open 'c:\program files\Microsoft Identity Integration Server\data' folder and note that the IIFP database (.MDF) and log file (.LDF)

are in this folder.

45. Create the same folder structure on the D: (logs) and the E: (database) volumes.

46. Copy (not move) the LDF to the '**d:\Program Files\Microsoft Identity Integration Server\data**' folder and copy (not move) the MDF to the '**e:\Program Files\Microsoft Identity Integration Server\data**' folder.

47. Switch back to the SQL Management Studio and select the 'Databases' node, right click and select 'Attach'.

48. When the 'Attach Database' dialog pops up, click the [Add] button

49. Select the 'MicrosoftIdentityIntegrationServer' database file from the *NEW* location (E:) and then click [OK].

50. Notice that once the MDF file loads, the current file path for the logs is no longer correct. Click the small 'ellipses' button next to the log file path (incorrectly showing as C:) and select the appropriate log file and click [OK]

51. Once both of the paths show as correct in the 'Attach Database' dialog box, click [OK]

52. An 'error' dialog will likely pop up; click [OK]

53. If you go back and click the hyperlink in the message column, you will be presented with another dialog that explains why the error occurred:

This is benign (and expected) so click [OK]. The database did attach correctly.

54. Click [Cancel] to dismiss the 'Attach Databases' dialog

55. Select the 'Databases' node and press [F5] to refresh the listing. You should see the 'MicrosoftIdentityIntegrationServer' database appear in the listing.

56. Open the properties for the database, then select the 'Files' page, and scroll to the right to verify that the database file and logs are being referenced in

the proper location.

57. Scroll the window back to the point where the 'Autogrowth' column is visible. Notice that the transaction log is configured with a 2GB size limit. Click the 'ellipses' button to bring up the autogrowth dialog and change the 'Maximum File Size' to be 'Unrestricted File Growth' and click [OK]

58. Click [OK] once more to dismiss the database properties dialog.

59. Switch back to the 'Computer Management' console and start the 'Microsoft Identity Integration Server' service. Once it starts, stop it again.

60. Switch back to explorer and create the following directory structure on drive F:

61. \Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\Data

62. Switch back to the SQL Management Studio and click the 'New Query' button

63. Copy and paste the following into the query pane:

USE master;
GO
ALTER DATABASE tempdb
MODIFY FILE (NAME = tempdev, FILENAME = 'F:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\data\tempdb.mdf');
GO
ALTER DATABASE tempdb
MODIFY FILE (NAME = templog, FILENAME = ' F:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\data\templog.ldf');
GO

64. Click Execute. Once the query executes, the screen should indicate the

following:

65. Switch back to the 'Computer Management' console and restart the 'SQL Server (MSSQLSERVER)' service.

66. Restart the 'Microsoft Identity Integration Service'

67. Go back to SQL Server Management Studio and open a new query pane.

68. Copy and paste the following to verify that tempdb has moved to the new location:

SELECT name, physical_name
FROM sys.master_files
WHERE database_id = DB_ID('tempdb');

69. Click Execute.

70. In the physical_name column, you should see the path to the new location:

71. You should also note the new tempdb files in the proper location in the file system, as well.

IIFP Management Agent Configuration

The configuration of the management agents for the GALSync server is documented below. In order to synchronize your e-mail addresses with the GALSync forest and other agencies, you will need to configure two management agents:

- One for your production forest
- One for the GALSync forest

Each management agent (MA) is then responsible for the import/export of data to its respective destination.

Note

Perform the following configuration logged on to the IIFP server as your production forest GALSync service account.

Production Active Directory forest MA

1. Open the Identity Manager

2. Select the 'Management Agent' toolbar button to switch to the MA context.

3. Go to the Actions menu and select Create.

4. On the Create Management Agent screen, specify a Management Agent for – and select 'Active Directory Global Address List (GAL)' from the drop-down combo-box.

Specify a name for the Management Agent, and then click [Next].

Note

For this management agent, use the agency abbreviation. I.E., KDE for the Kentucky Department of Education, LRC for Legislative Research Commission, COT for the Commonwealth Office of Technology, KRS for Kentucky Retirement Systems, etc.

5. On the **'Connect to Active Directory Forest'** page, enter in the root DNS name of your production forest, along with the GALSync service account (created by you for your forest), its password and the DNS name of the root domain

and then click [Next]

On the 'Configure Directory Partitions screen', select the Directory Partitions corresponding to the domains that have users that you wish to publish to the GALSync forest and then click the [containers] button to select the OU's within those domains. You will likely find it easier to unselect the root domain object in the containers dialog and then just selectively select what you wish to synchronize.

See the diagram on the right for an example:

Note

6. On the **'Configure GAL screen'**, navigate to the 'GAL Container Configuration frame and click the [Target] button and select the container that will hold the contacts that are imported from the GALSync forest (which represent every state user except for your agency).

    Click the [Source] button and for each directory partition that you have selected the IIFP to import from (refer to the previous step) specify which OU's contain users that you wish to synchronize to the GALSync forest (for the purpose of every other state agency being able to import into their GAL). Note that if you have multiple domains and multiple OU's to synchronize that this step will likely be fairly lengthy and/or tedious.

Once finished adding all of the source containers, then proceed down to the 'Exchange Configuration' frame and click the [edit] button to add the listing of authoritative SMTP domains that are in your forest, as shown below:

click [OK] to accept all of the SMTP domains you have entered into the dialog box, and then click [Next] to proceed to the next step. Leave all other settings at their default for this page!

7. On the 'Select Object Types' page, leave all options at their defaults and click [Next]

8. On the 'Configure Connector Filter' page, leave all options at their defaults and click [Next]

9. On the 'Configure Join and Project Rules' page, select 'group' from the 'Data Source Object Type' column, then select its project rule from the 'Join and projection rules for: group' frame and click [Delete].

After making this configuration change, the screen should look like the following:

Click [Next] to proceed.

10. On the 'Configure Attribute Flow' page, expand the Object Type: user to Object Type: person. Click the 'ProxyAddresses' <- 'legacyExchangeDN' attribute flow so that it is highlighted.

Click [Delete] to remove the attribute flow mapping, and then click [Next] to continue.

11. On the 'Configure Deprovisioning' page, accept all defaults, and click [Next] to continue.

12. On the 'Configure Extensions' page, accept all defaults and click [Finish] to complete configuration of the MA and allow the IIFP to create the MA with these options.

COK Galsync Forest MA

1. Go to the Actions menu and select Create.

2. On the Create Management Agent screen, specify a Management Agent for – and select 'Active Directory Global Address List (GAL)' from the drop-down combo-box.

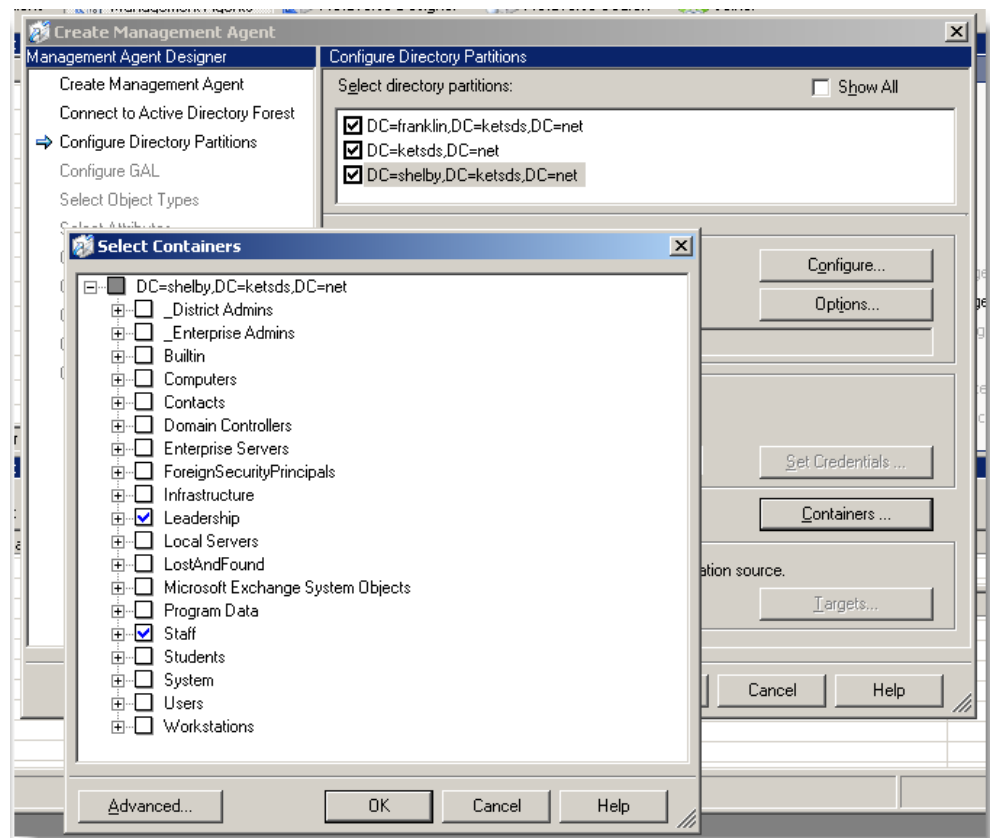Specify the name for the Management Agent to be 'GALSync Forest' and then click [Next].

3. On the **'Connect to Active Directory Forest'** page, enter in the root DNS name of your production forest, along with the GALSync service account (created by you for your forest), its password and the DNS name of the root domain

and then click [Next]

4. On the 'Configure Directory Partitions' page, place a checkbox in the 'DC=GALSync,DC=ky,DC=gov' partition, then click the [Containers] button to bring up the following dialog to select the containers that you wish to receive contacts from (corresponding to the agencies that you wish to see in your GAL).   Be sure to select your own container in this screen as well. We will differentiate between the OU's next. Click on [OK] to close this dialog when finished and then click on [Next] button to continue.

5. On the 'Configure GAL' page click the [Target] button to select the OU that your agency 'owns' in the GALSync forest. This container will be named after your agency and will be the OU that all of your mail and mailbox-enabled users will be written to the GALSync forest as contacts.

   After selecting your synchronization OU, click the [Source] button to select the containers that hold the contacts that you wish to show up in your GAL. You have total control as to which containers you select here, and you are not required to synchronize an agency if you do not want to see their users in your GAL. Click [OK] when finished.

6. Click the [Edit] button in the 'Exchange Configuration' frame and enter in all of the SMTP suffixes that your forest is authoritative for, as shown below:

Click [OK] when finished and then click [Next] to proceed.

7. On the 'Select Object Types' page, leave all options at their defaults and click [Next]

8. On the 'Configure Connector Filter' page, leave all options at their defaults and click [Next]

9. On the 'Configure Join and Project Rules' page, select 'group' from the 'Data Source Object Type' column, then select its project rule from the 'Join and projection rules for: group' frame and click [Delete].
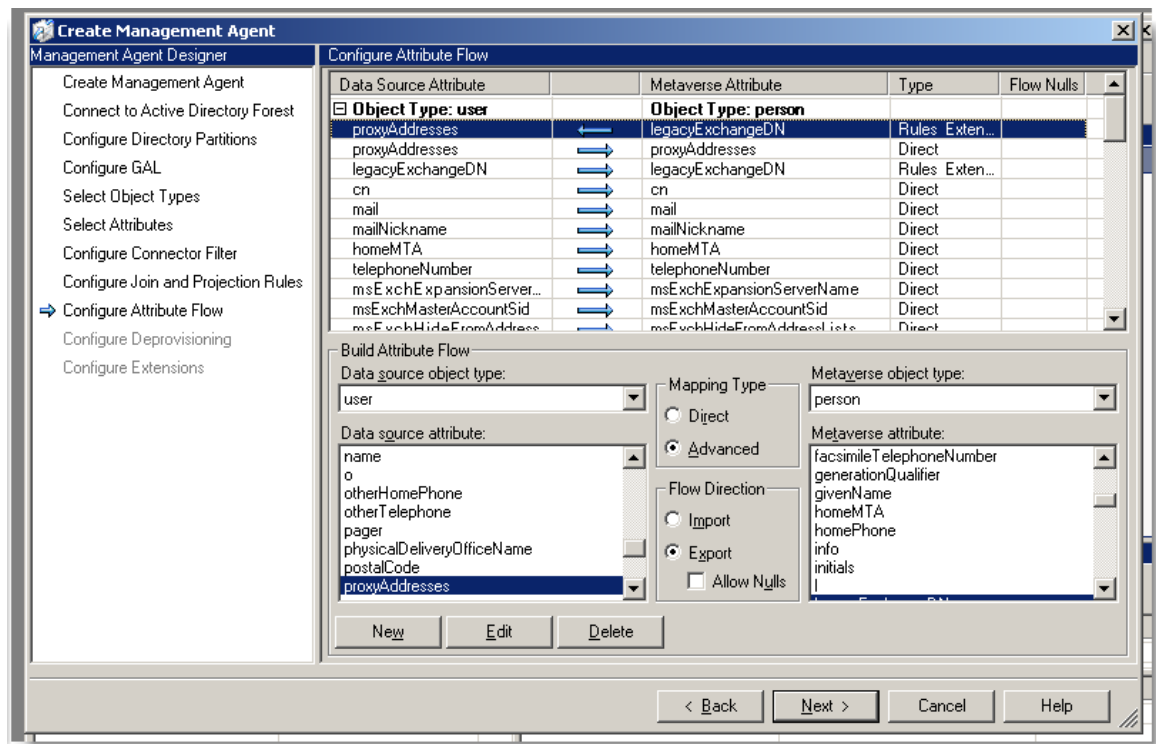

After making this configuration change, the screen should look like the following:

Click [Next] to proceed.

10. On the 'Configure Attribute Flow' page, expand the Object Type: contact to Object Type: cot_galsync_forest. Click the 'ProxyAddresses' <– 'legacyExchangeDN' attribute flow so that it is highlighted.
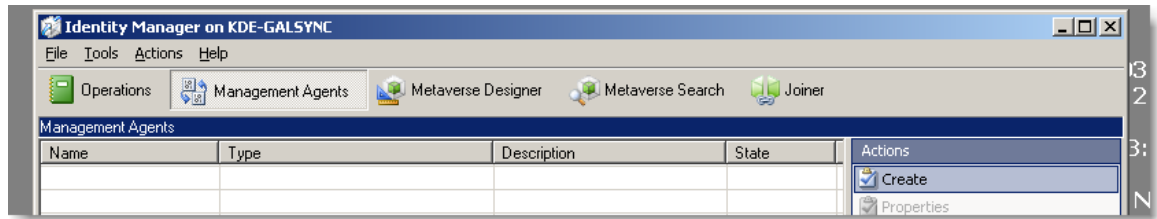

Click [Delete] to remove the attribute flow mapping, and then click [Next] to continue.
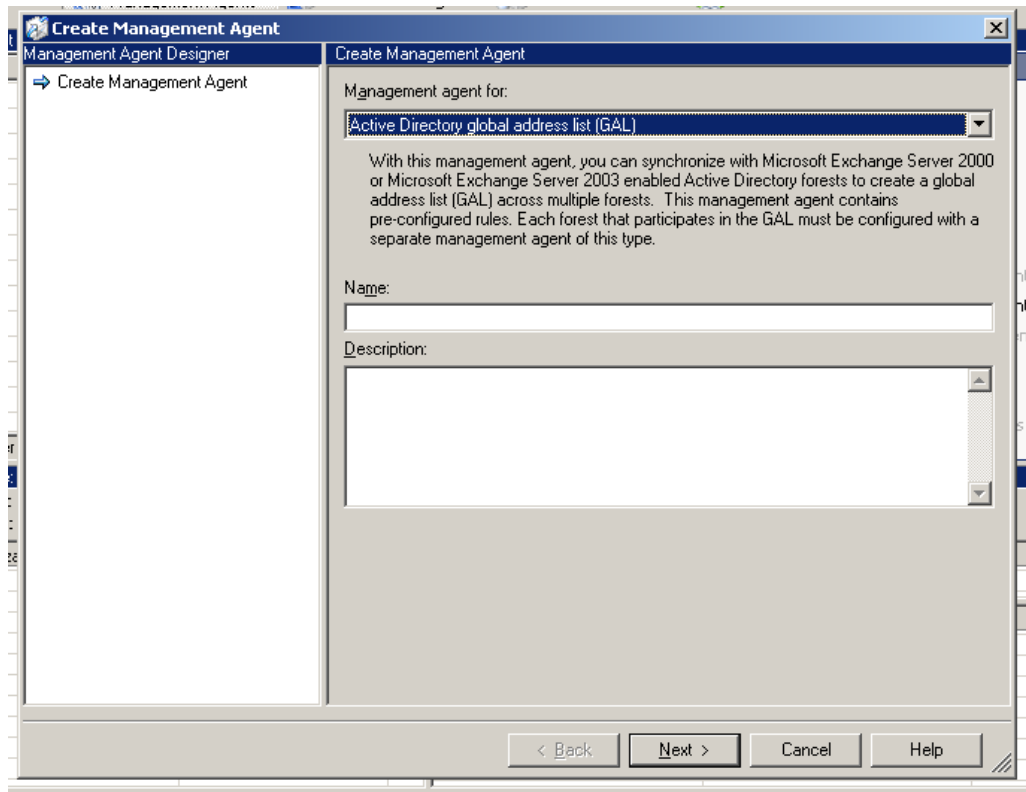
13. On the 'Configure Deprovisioning' page, accept all defaults, and click [Next] to continue.

11. On the 'Configure Extensions' page, accept all defaults and click [Finish] to complete configuration of the MA and allow the IIFP to create the MA with these options.

IIFP Server General Configuration Options

Once the MA configurations have been completed, go back to the main Identity Manager window, and select the Tools/Options menu and turn on the provisioning rules extensions (with the checkbox shown below) and click [OK] to save the changes.

If this is not enabled, then provisioning (creation of new contacts) will not run and although synchronization will appear to work correctly, nothing will be created in either forest.

IIFP Extension Code Update

Browse to the 'IIFP\Extension Code' folder on the 'COK GALSync Resources' CD. Notice the single 'GALSync.dll' file in that folder. Copy that file to the 'C:\Program Files\Microsoft Identity Integration Server\Extensions' folder and overwrite the existing GALSync.dll file (make sure to keep a backup of the 'as-installed' file!!)

IIFP Attribute Synchronization Configuration

The default configuration of the management agents (as has been configured thus far) will synchronize the following attributes from your production Active Directory forest to the GALSync forest and from the other agencies into your production Active Directory forest (as contacts):

Listing 1 – Primary Attributes

1. company
2. department
3. displayName
4. givenName (first name)
5. l (city)
6. mailNickname (mail alias)
7. physicalDeliveryOfficeName (office)
8. postalCode (zipcode)
9. Sn (last name)
10. St (state)
11. streetAddress
12. targetAddress (e-mail address for contact)
13. telephoneNumber
14. title

Listing 2 – Secondary Attributes

1. cname
2. groupType
3. hideDLMembership
4. homeMDB
5. homeMTA
6. legacyExchangeDN
7. mail
8. mAPIRecipient
9. msExchExpansionServerName
10. msExchHomeServerName
11. msExchMasterAccountSid
12. msExchOriginatingForest
13. msExchPoliciesExcluded
14. msExchHideFromAddressLists
15. nTGroupMembers
16. ProxyAddresses

Listing 3 – Ancillary Attributes

1. c
2. co
3. division
4. employeeID

5. employeeType
6. facsimilieTelephoneNumber
7. homePhone
8. info
9. initials
10. manager
11. mobile
12. msExchAssistantName
13.
14. name
15. o
16. otherHomePhone
17. otherTelephone
18. pager
19. telephoneAssistant
20. userCertificate
21. userSMIMECertificate

The attributes listed above are broken down as follows:

1. The attributes in **'Listing 1'** are the 14 agreed-upon attributes that will be exchanged as part of the GALSync solution.
2. The attributes in **'Listing 2'** are attributes that are 'internal' to the GALSync process and are required in order for the process to work correctly, but will never show up in a participating agency's GAL, so the information in these attributes is deemed 'safe' to synchronize and doesn't compromise any participating entity in any way, shape or form.
3. The attributes in **Listing 3** are attributes that are not part of the 14 agreed-upon attributes and are not essential to the GALSync solution working correctly. Under the default configuration, whatever data is in these attributes will be synchronized to the GALSync forest and will be pulled from the GALSync forest and populated on the synchronized contacts from the other agencies and could potentially show up in your GAL.

Note

If either of these scenarios present a data protection risk for your agency, then you are advised to 'turn off' the offending attributes and choose to either not push them to the GALSync forest for other agencies to have in their GAL and/or not pull them from the GALSync forest and publish them to your GAL. Since each participating agency owns its GALSync server, it is the participating agency's responsibility to properly configure the IIFP to only synchronize what is deemed 'safe'.

COT bears no responsibility for an agency accidentally publishing inappropriate data and bears no responsibility for an agency accidentally synchronizing inappropriate data into their forest.

In order to configure your management agents to not synchronize a specific attribute, perform the following:

In this example, we will turn off the replication of the 'employeeID' attribute. Disabling replicate of any other attribute will follow the same principle and configuration steps.

Note

In order to control what is synchronized from your production Active Directory forest to the GALSync forest you must modify the MA for your Active Directory forest. In order to control what is synchronized from other agencies into your GAL, you must modify the 'GALSync Forest' MA.

---

1. Open the properties of the management agent in question, and go to the 'Select Attributes' page:

2. Select the 'employeeID' attribute and uncheck the checkbox next to the name to deselect it from being managed by this MA.

3. Now select the 'Configure Attribute Flow' page and click [Delete] to remove the attribute mapping where desired. In order to completely negate the flow of this attribute, it must be removed as an import flow from the 'user' to 'person' object mapping, the export flow from the 'contact' to 'contact_galsync_forest' object mapping, the export flow from the 'contact' to 'person' object mapping and the import flow from the 'contact' to 'contact_cot' object mapping, as shown below:

4. Once finished, click [OK] to save the MA configuration

5. Repeat the above process for any other attribute on the 'Ancillary Attributes' list that you wish to remove from your configuration.

Active Directory

This section details the configuration changes that will need to be made to Active Directory in order to support location of the GALSync forest domain controllers and authentication to the servers.

DNS

Perform the following on a DNS server (domain controller) in the domain where the GALSync server is installed:

1. Start the DNS manager
2. Expand the DNS tree down to the 'Forward Lookup Zones' node
3. Right-click and select 'New Zone…' from the context menu
4. Click [Next] to begin the new zone creation wizard
5. Accept the option to create a new forward lookup zone, but uncheck the option to store the data in Active Directory (at this time).

Click [Next] to proceed.

6. Name the zone 'GALSync.ky.gov' and click [Next]

Click [Next] to proceed.

7. Select the option for using the existing 'galsyn.ky.gov' zone file.   Before proceeding to the next step, open the 'DNS' folder on the 'COK GALSync Resources' CD and copy the 'GALSync.ky.gov.dns' file from the CD to the 'C:\windows\system32\dns' folder on the DNS server.

Click [Next] to proceed.

8. Verify that dynamic updates are turned off

Click [Next] to proceed.

9. Click [Finish] to complete the Zone Creation wizard.

10. Once the zone is created and loaded from the file copied from the CD, you should be able to click on the zone in the DNS manager to see the data that was imported from the text file. Assuming that the import worked as expected, right-click on the zone and open its properties dialog:

On the 'General' tab click the [Change] button to bring up the dialog to allow the change from a standard (file-based) to an Active Directory-integrated zone:

11. Select the checkbox to turn on Active Directory storage for the zone data, and click [OK}

12. Verify that the data is now stored in Active Directory

and click [OK] to save the changes to the zone.

Exchange 2003

Recipient Policy

If your production Exchange 2003 environment is configured with multiple recipient policies that use specific LDAP filters to evaluate different objects, then ensure that you have a recipient policy that will apply to the synchronized contacts. By default, the synchronized contacts will be evaluated by the 'Default Policy' (which is the lowest-valued policy). If your default policy will not evaluate the synchronized contacts (due to visibility requirements) and no other policy will evaluate the synchronized contacts, then be prepared to create a special recipient policy for the synchronized contact objects.

If this is the case, then modification of the attribute flow rules on the IIFP server will also be necessary, so that the IIFP is configured to flow a specific attribute/value pair (of your designation) so that the synchronized contacts will meet at least one recipient policy filter condition and their 'showInAddressBook' attribute will be updated accordingly. Note that this information will not be flowed back to the GALSync forest – this is done only to ensure that the contacts work within your environment.

Note

Failure of the RUS to update the 'showInAddressBook' attribute on the synchronized contacts will prevent the synchronized contacts from being visible in any Exchange 2003 address list (Global, Offline, etc.)

Recipient Update Service (RUS)

If your production Active Directory is configured in a multi-domain environment then ensure that you have a recipient update service (RUS) configured appropriately within the Exchange System Manager (ESM) to update the synchronized contact objects from the GALSync forest. Although the contacts themselves will not have their e-mail address information updated, the RUS is required to evaluate recipient policy information so that the 'showInAddressBook' attribute is populated correctly on the contact objects such that they show up correctly in the appropriate address lists.

Backup and Recovery of the GALSync Server

This section details the high-level process by which to backup and restore the GALSync server.

General Information

Since the GALSync server is really a client/server application, there are a several components that must be accounted for to completely back up the GALSync server's configuration:

- SQL databases
- IIFP MA configuration
- IIFP encryption keys
- Extension code

Each of these components comprise a portion of the total GALSync server configuration and all must be backed up in order to return the server to its production configuration after a disaster.

This section covers these topics at a high level; much more detailed information can be found in the Microsoft **"Maintaining the MIIS 2003 Database"** whitepaper, which is located at the following location on the Microsoft website - http://www.microsoft.com/downloads/details.aspx?FamilyID=5f50fea4-b876-49ef-a9b9-593e48b5d640&DisplayLang=en and is also located on the "COK GALSync Resources" CD ISO file in the 'Documents' Folder, along with a copy of this document file.

Although the information in this section will allow you to perform an initial backup and a restoration of your server, please review the aforementioned Microsoft whitepaper, as the long-term care and maintenance of the SQL database is covered in this document and is outside the scope of this implementation.

Backup

The SQL 2005 server that runs on the IIFP server should be configured with a maintenance plan to automatically optimize and backup the SQL database to a spindle on the IIFP server and then secure the backup to some location on the network or take it to tape for archival purposes.

The information in the IIFP database is encrypted and can only be decrypted with the server's encryption key, so restoring the database to another location or server without the key will not be fruitful.

Also be sure to export the Management Agents to XML files and then keep these XML files in the same location as the encryption key. Exporting a MA can be performed as follows:

1. Open the Identity Manager and switch to the 'Management Agents' context.
2. Either right-click on the MA that you wish to export and select 'Export Management Agent' from the context menu, or select 'Export Management Agent' from the Actions menu on the right-hand side of the screen, as shown below:

3. The normal Windows 'Save As' dialog will then appear; name the export XML file appropriately and then click [Save] to begin the exporting process.

Backup of the encryption key can be performed as follows:

1. Open the 'c:\program files\microsoft identity integration server\bin' folder and execute the 'miiskmu.exe' file

2. Click [Next]

3. Specify the IIFP service account credentials (and NETBIOS domain name!) and click [Next]

Specify a location and name for the exported key set file

4. Click [Finish] to begin the exporting process

5. Click [Close] to exit the wizard and open the folder where the exported encryption key file was created, and move the key file to a safe location!

Restoration

If your GALSync server befalls some type of disaster, or it becomes necessary to move it to another server, then the server will have to be restored from its last good backup in order to return it to service and 'catch-up' on the synchronization process.

The full restoration process is as follows (all procedures must be performed logged on as the GALSync service account):

1. Bring up a new server that meets the specifications as outlined in this document
2. Install SQL 2005 standard edition
3. Install SQL 2005 SP2
4. Restore the last backup of the 'MicrosoftIdentityIntegrationServer' database that you have.
5. Install the IIFP from the 'COK GALSync Resources' CD
6. During the installation, the setup program will detect that an existing MIIS/IIFP database is on the server and will prompt you as to whether or not you want to restore the configuration associated with that database. Select [Yes] to do so.
7. Run the 'miisactivate' utility to restore the encryption key set
8. Log off/on to the server
9. Install the 'KB893015' patch to the IIFP server from the 'COK GALSync Resources' CD
10. Ensure the MIIS services are started and go into the Identity Manager and verify that all prior information (run history, MA's, etc.) are listed. If the restore process completed successfully, everything in its entirety should be there and the server should be ready to continue service.
11. Update the 'GALSync.dll' file from the 'COK GALSync Resources' CD and copy it over the existing 'GALSync.dll' in the 'c:\program files\microsoft identity integration server\extensions' folder
12. Ensure that provisioning is turned on in the server options.

Automating the Synchronization Process

Since the GALSync server runs at each local agency, there is no 'centralized' GALSync process that has to be accounted for in order for the entire procedure to work correctly. As such, each agency is free to determine what import and export schedule (frequency, time of day, day of week, etc.) works best for them.

As with any synchronization solution, the value is provided in the automation of the process and the currency of the data. As such, COT advises synchronize on a semi-regular basis (at least once per week) if for no other reason than to provide the rest of the Commonwealth with your updated information.

COT recommends that you time your exports during the day such that the exported information has replicated throughout your environment *BEFORE* the OAB (Offline Address Book) files are rebuilt on the Exchange server. This recommendation is done such that users download the most recent OAB changes every morning and their 'view' of the GAL is as up-to-date as possible.

If you are following a 'once-a-week' schedule, then the weekend (Sunday during the day) is a perfect time for the exports. If you are following a daily schedule, then time the exports such that they complete before the OAB rebuild for that day begins.

The IIFP supports running the MA's manually at your discretion. However, if you wish to automate the running of the MA's you will need to install the MIIS Resource Kit utilities. You can download the newest version of the utilities from the Microsoft ILM website, or just install them directly from the 'IIFP\Resource Kit' folder on the 'COK GALSync Resources CD'.

Once the utilities are installed, open the installation folder (C:\Program Files\Microsoft Identity Integration Server 2003 Resource Tool Kit\) and find the 'masequenceconfiguration.exe' file. Run this utility and from here the sequence of the management agents can be configured. Once the sequence is configured correctly, then save the sequence as an XML file and run the command-line program 'masequencer.exe' and use the XML file as input to run the MA's in the order you configured.

Once this configuration is solid, then create a scheduled task to run 'masequencer' with the final XML configuration file to fully automate the GAL Synchronization process.

Performing your first Full GALSync

As soon as you have your server set up and configured, you should be ready to perform your first GALSync. There are a few items to keep in mind and you should consider the following **before** jumping right into the process and performing your first synchronization.

- Plan on performing your first synchronization after-hours or on a weekend – the time required for the initial setup of the server and the preparation of Active Directory will likely necessitate some uninterrupted time to think through all of the processes and procedures.
- Ensure that you have your GALSync forest credentials from COT
- Ensure that you can resolve the 'gs1.galsync.ky.gov' and the 'gs2.galsync.ky.gov' servers from the GALSync server and that you can telnet to ports 389 and 88 to both servers.
- Prepare your users for the ramifications of the first synchronization cycle. Things to make them aware of consist, but are not limited to, the following:
  - The names of the state contacts they are used to seeing in the GAL for a specific user may change or be formatted slightly differently
  - Contacts and distribution lists in their personal address books will not be altered by this process, but may not work after the first synchronization. Each user should be prepared to update their own PAB with the new synchronized contacts if they wish to continue using PAB's.
  - Replying to a message from a state contact that was created before the first synchronization process *MAY* not work and *MAY* result in a NDR. Every effort has been put forth to keep this from happening but the possibility does exist. Once the user selects the same person from the GAL and forwards the old message to them or creates a new message to them, this will be the last time such an NDR is seen.
  - Distribution list membership in the GAL may be inconsistent for a few days and they should allow a day or two for full membership to be re-established.
- All existing mail-enabled state contacts that correspond to e-mail addresses outside your agency should be deleted from Active Directory and force replication between the domain controllers in the domain where these contacts were located.
- All existing distribution lists in the GAL that currently have mail-enabled state contacts will see their membership modified (and those contacts removed) once the existing contacts are deleted before the first sync is performed. Once the first sync and export finished and these contacts are re-created, they will have to be added back to the distribution lists in question.

Note

This 're-add' process will only have to be done for the initial sync. Distribution list membership is maintained from that point going forward.

Once you are ready to perform the synchronization, follow this procedure for running the management agents:

1. Active Directory forest MA – perform a full import (stage only) of all directory partitions that contain mailbox-enabled users in your forest
2. GALSync forest MA – perform a full import (stage only)
3. Active Directory forest MA – perform a full synchronization
4. GALSync forest MA – perform a full synchronization
5. Active Directory forest MA – perform an export
6. GALSync forest MA – perform an export

7. **This step is important – make sure the Exchange RUS runs against your contact objects to make sure they show up in the GAL!!**

8. Active Directory forest MA – perform a delta import (stage only) of the directory partition that contains the contacts
9. Active Directory forest MA – perform a delta synchronization
10. Active Directory forest MA – perform an export
11. Active Directory forest MA – perform a delta import (stage only) of the directory partition that contains the contacts
12. GALSync forest MA – perform a delta import (stage only)

The 'Day-toDay' GALSync process

Even though each agency is in control of how often they 'push' their data to the GALSync forest and how often they 'pull' the data about other agencies from the GALSync forest, the true value of such a solution is in the

Although your needs may vary, the following sequence of MA runs should be followed for 'day-to-day' running of the synchronization process:

1. Active Directory forest MA – perform a delta import (stage only) of each directory partition
2. Active Directory forest MA – perform a delta synchronization of each directory partition
3. GALSync forest MA – perform a delta import (stage only)
4. GALSync forest MA – perform a delta synchronization
5. Active Directory forest MA – perform an Export
6. GALSync forest MA – perform an export

7. **This step is important – make sure the Exchange RUS runs against your contact objects to make sure they show up in the GAL!!**

8. Active Directory forest MA – perform a delta import (stage only) of the directory partition that contains the contacts
9. Active Directory forest MA – perform a delta synchronization
10. Active Directory forest MA – perform an Export
11. Active Directory forest MA – perform a delta import (stage only) of the directory partition that contains the contacts

Note

Be sure that your exports are timed such that the domain controllers in your forest have enough time to replicate the new contacts and before the Exchange 2003 server rebuilds its Offline Address Book so that all new contacts synchronized on a given day are available to all of your Outlook users (both online and cached-mode) the next day!

Appendix A – Process for Requesting and Establishing GALSync

History

During the Exchange 5.5 era agencies using the Kentucky Information Highway (KIH) could exchange the Global Address List (GAL) with the Executive Branch. This process was performed manually, tedious and time consuming. The GAL swap process was an official means for agencies not hosted on the Executive Branch's email system an avenue to update their system to show contacts from other agencies. All information included in the GAL is the property of and owned by the agencies.

Present Day

The installation and upgrade to Exchange 2003 has put an end to the manual process previously used to update the GAL. We have an opportunity to take a once manual process and turn it into an automated process for all areas that qualify to participate.

Method

The automation of swapping GALs will be accomplished by using Identity Integration Feature Pack (IIFP) which is the scaled-down version of Microsoft Identify Integration Server (MIIS). IIFP provides all of the functionality to accomplish the needed business goals while providing an tremendous licensing savings to each participating agency (the IIFP is free). The agency IIFP will run an automated synchronization cycle on a weekly basis to perform this synchronization with the GALSync forest.

The Executive Branch's Commonwealth Office of Technology will provide and maintain the staging (GALSync) forest, which is compiled of mail-enabled contacts from the participating agencies. It will be up to each participating agency to provide their own IIFP/SQL server and configuration of this server that will provide the correct information to the GALSync forest.

Requirements

The qualifications to participate in the GAL Exchange via IIFP are as follows:
- An Agency's GAL and its published data is their property and the use of it by someone else needs to be approved by the proper management of that Agency.

- An agency needs to be part of the Commonwealth's network which has been referred to as "The Kentucky Information Highway" per KRS 45A.605.

    COT Responsibilities - COT will be responsible for the implementation of the first IIFP server and staging forest that all participants will use to house their contacts (and will incur any hardware, software and ongoing maintenance/monitoring costs). COT will maintain the hardware and software to insure proper functionality of the staging forest server and will make every effort to ensure that the forest is available at all times. If there is any scheduled downtime, COT will notify you (sent to the agency contact e-mail address submitted on the Acknowledgement Letter) of any downtime as far in advance as possible.

Participants Responsibilities - All participants understand they are responsible for acquiring their own IIFP server and any configuration that will be required to achieve the exchanging of GAL contacts. The agencies will be responsible for any support required to implement this solution as well as any ongoing support that may be required for the proper functionality of their IIFP server.

Required Fields for all GALSync Contacts

By participating in the GALSync program, your agency agrees to publish the following information to all other participating agencies. By participating in the GALSync program, your agency will automatically 'learn' the same information about all other participating agencies, as well. The Active Directory attribute names are given in red for your convenience.

Synchronized User Information	Active Directory Attribute Name
First Name	givenName
Last Name	sn
Display Name	displayName
E-mail Alias	mailNickname
E-mail Address	mail
Address	streetAddress
City	l
State	st
Zip Code	postalCode
Title	title
Company	company
Department	department
Office	physicalDeliveryOfficeName
Phone	telephoneNumber

The Commonwealth Office of Technology will require a Letter of Acknowledgement (see next page) that all of the above requirements will be met and abided by.

Commonwealth of Kentucky GALSync Authorization and Letter of Acknowledgement

By completing and submitting this letter to the Commonwealth Office of Technology, you hereby declare the following:

The agency which I represent wishes to participate in the COT-sponsored 'Global Address List Synchronization Program' - hereafter referred to as 'GALSync'. By participating in this program, this agency hereby releases COT from any responsibility for such data that is synchronized and any adverse affect that is created or experienced by following the instructions included in this document.

COT has tested – to the best of its ability – all instructions, procedures, events and outcomes noted within, but cannot be held responsible for unforeseen circumstances and configurations within your agency. The participating agency has the responsibility to prepare their environment for this program, and by requesting to participate in such program, you deem yourself capable and willing to follow these directions (in their entirety) and to be able to recover your environment in the event of an unsuccessful implementation.

In order to participate in the GALSync program, COT requires the following information from your agency:

Description	Value
Agency Name	
Agency Contact Name	
Agency Contact E-mail address	
Dedicated 'Public' IP address from Agency firewall for the IIFP server	

Once this information has been submitted to COT, then you will receive back a copy of this form along with your agency-specific credentials and connectivity information needed to access the GALSync forest to perform GALSync.

These credentials will **only allow your agency to update your agency's information into the GALSync forest**, and to **read the information from all other participating agencies**. Under these conditions, it is very important to ensure that the attributes that will be published to the other participants GAL contains only the information that your agency deems non-sensitive and has the legal authority to share with other participating agencies.

Agency Representative
   Signed By: _____

   This _____ day of _____, _____

COT Representative
   Signed By: _____

   This _____ day of _____, _____

Appendix B - How to grant the "Replicating Directory Changes" permission for the Microsoft Metadirectory Services ADMA service account

[View products that this article applies to.](#)

 Article ID   : 303972
 Last Review : February 23, 2007
 Revision   : 5.1

This article was previously published under Q303972

On This Page

↓SUMMARY
↓MORE INFORMATION
↓Setting permissions by using the ACL editor
↓Setting permissions by using Adsiedit

SUMMARY

When discovering objects in Active Directory using the Active Directory management agent (ADMA), the account that is specified for connecting to Active Directory must either have Domain Administrative permissions, belong to the Domain Administrators group, or be explicitly granted Replicating Directory Changes permissions for every domain of the forest that this management agent accesses. This article describes how to explicitly a grant a user account the Replicating Directory Changes permissions on a domain.

Note In Windows Server 2003, the name of this permission changed to "Replicate Directory Changes."

⇑Back to the top

MORE INFORMATION

The Replicating Directory Changes permission, known as the Replicate Directory Changes permission in Windows Server 2003, is an Access Control Entry (ACE) on each domain naming context. You can assign this permission by using the ACL editor or the Adsiedit support tool in Windows 2000.

⇑Back to the top

Setting permissions by using the ACL editor

1. Open the Active Directory Users and Computers snap-in

2. On the **View** menu, click **Advanced Features**.

3. Right-click the domain object, such as "company.com", and then click **Properties**.

4. On the **Security** tab, if the desired user account is not listed, click **Add**; if the desired user account is listed, proceed to step 7.

5. In the **Select Users, Computers, or Groups** dialog box, select the desired user account, and then click **Add**.

6. Click **OK** to return to the **Properties** dialog box.

7. Click the desired user account.

8. Click to select the **Replicating Directory Changes** check box from the list.

9. Click **Apply**, and then click **OK**.

10. Close the snap-in.

⇑Back to the top

Setting permissions by using Adsiedit

Warning Using Adsiedit incorrectly can cause serious problems that may require you to reinstall your operating system. Microsoft cannot guarantee that problems resulting from the incorrect use of Adsiedit can be solved. Use Adsiedit at your own risk.

1. Install the Windows 2000 Support tools if they have not already been installed. For more information about how to install the Windows 2000 Support Tools, click the following article number to view the article in the Microsoft Knowledge Base:

    301423 How to install the Windows 2000 Support Tools to a Windows 2000 Server-based computer

2. Run Adsiedit.msc as an administrator of the domain. Expand the Domain Naming Context (Domain NC) node. This node contains an object that begins with "DC=" and reflects the correct domain name. Right-click this object, and then click **Properties**.

3. Click the **Security** tab.

4. If the desired user account is not listed, click **Add**, otherwise proceed to step 8.

5. In the **Select Users, Computers, or Groups** dialog box, select the desired user account, and then click **Add**.

6. Click **OK** to return to the **Properties** dialog box.

7. Click **Apply**, and then click **OK**.

8. Select the desired user account

9. Click to select the **Replicating Directory Changes** check box.

10. Click **Apply**, and then click **OK**.

11. Close the snap-in.

Note Using either method, setting the Replicating Directory Changes permission for each domain within your forest enables the discovery of objects in the domain within the Active Directory forest. However, enabling discovery of the connected directory does not imply that other operations can be performed.

To create, modify, and delete objects within Active Directory using a non-administrative account, you may need to add additional permissions as appropriate. For example, for Microsoft Metadirectory Services (MMS) to create new user objects in an Organizational Unit (OU) or container, the account that is being used must be explicitly granted the Create All Child Objects permission, as the Replicating Directory Changes permission is not sufficient to allow the creation of objects.

In a similar fashion, the deletion of objects requires the Delete All Child Objects permission.

It is possible that there are limitations on other operations, such as attribute flow, depending on the specific security settings that are assigned to the object in question, and whether or not inheritance is a factor.

⇧Back to the top

APPLIES TO

-   Microsoft Metadirectory Services 2.2 Service Pack 1
-   Microsoft Metadirectory Services 2.2 Service Pack 1
-   Identity Integration Feature Pack for Microsoft Windows Server Active Directory
-   Microsoft Identity Integration Server 2003 Enterprise Edition

⇧Back to the top